

The Packing of Spheres

What is the densest way to arrange identical spheres in space? There has been much progress on the problem, particularly in 24 dimensions, and the results can be applied to digital signaling

by N. J. A. Sloane

A manufacturer of ball bearings is asked to deliver as many balls, all the same size, to a foreign port as he can ship on a day's notice. Although the ball bearings are ready, only one ship is available, and the shipper explains to the manufacturer that the draft of the ship would be too deep for the local channel if the balls were to fill more than three-fourths of the volume of the hold. The manufacturer is unperturbed. "Your ship is safe," he replies. "Let the hold be filled to the hatches." Should the shipper believe his client?

The solution to this problem depends on finding out how densely a large number of identical spheres can be packed together in space. If, instead of ball bearings, the ship's hold were to be packed with cubes all the same size, the answer would be easy. Since the cubes would fit together with no wasted space in between, the hold could be essentially filled with cubes (ignoring the small spaces that might be left around the walls and ceiling) and the manufacturer's assurances would clearly be wrong. Balls, however, cannot be packed without wasting space. If in spite of arranging the ball bearings as densely as possible the wasted space still exceeds a fourth of the volume of the hold, the shipper can safely fill the hold with balls and proceed out of the channel.

A few minutes spent experimenting with oranges or billiard balls is enough to mislead many people into thinking the problem is trivial. Arrange three spheres on a flat surface so that their centers form an equilateral triangle; continue adding spheres on the surface so that each new sphere touches at least two spheres already in place. Build a second layer of spheres by placing each new sphere in the "deep hole," or depression, left at the center of any triangular group of spheres in the first layer. The finished second layer is identical with the first layer, although it is shifted in the horizontal plane. If more layers are added in the same way, the packing of spheres that results is called the

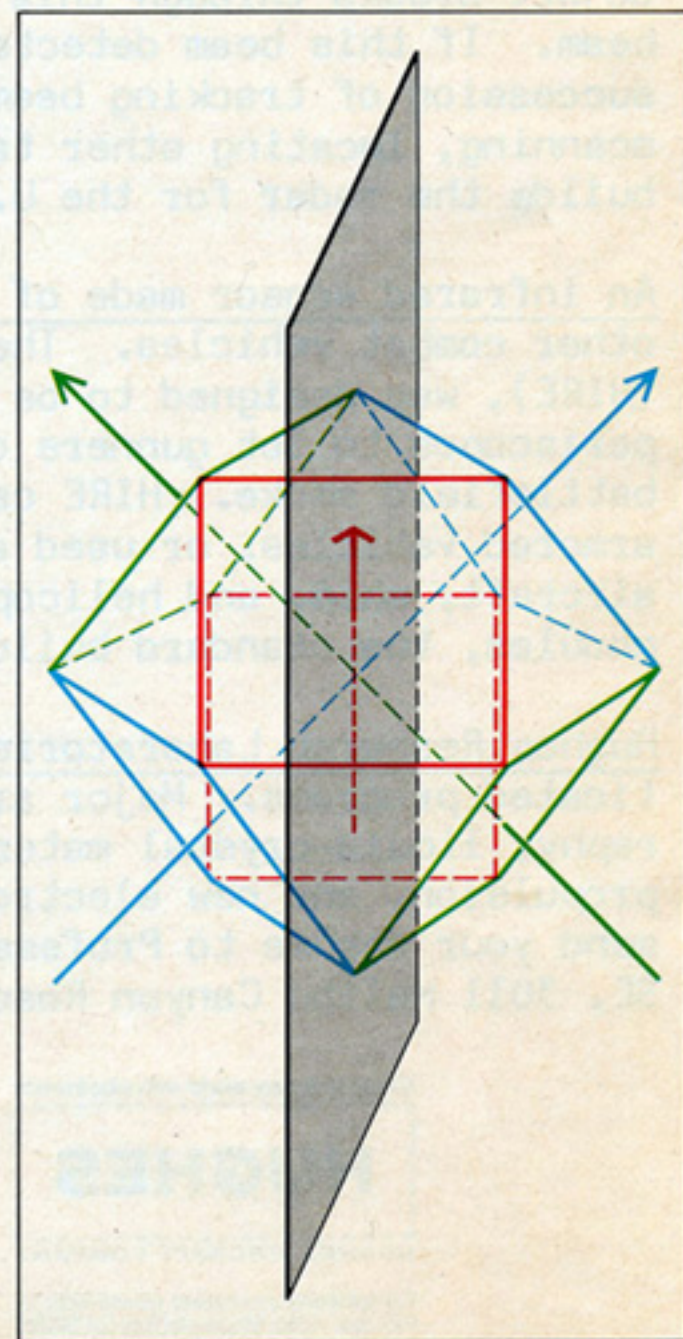
face-centered-cubic packing, which is familiar to chemists and crystallographers; it fills just over 74 percent of the volume of the space. As far as anyone knows, it is the densest packing that can be achieved.

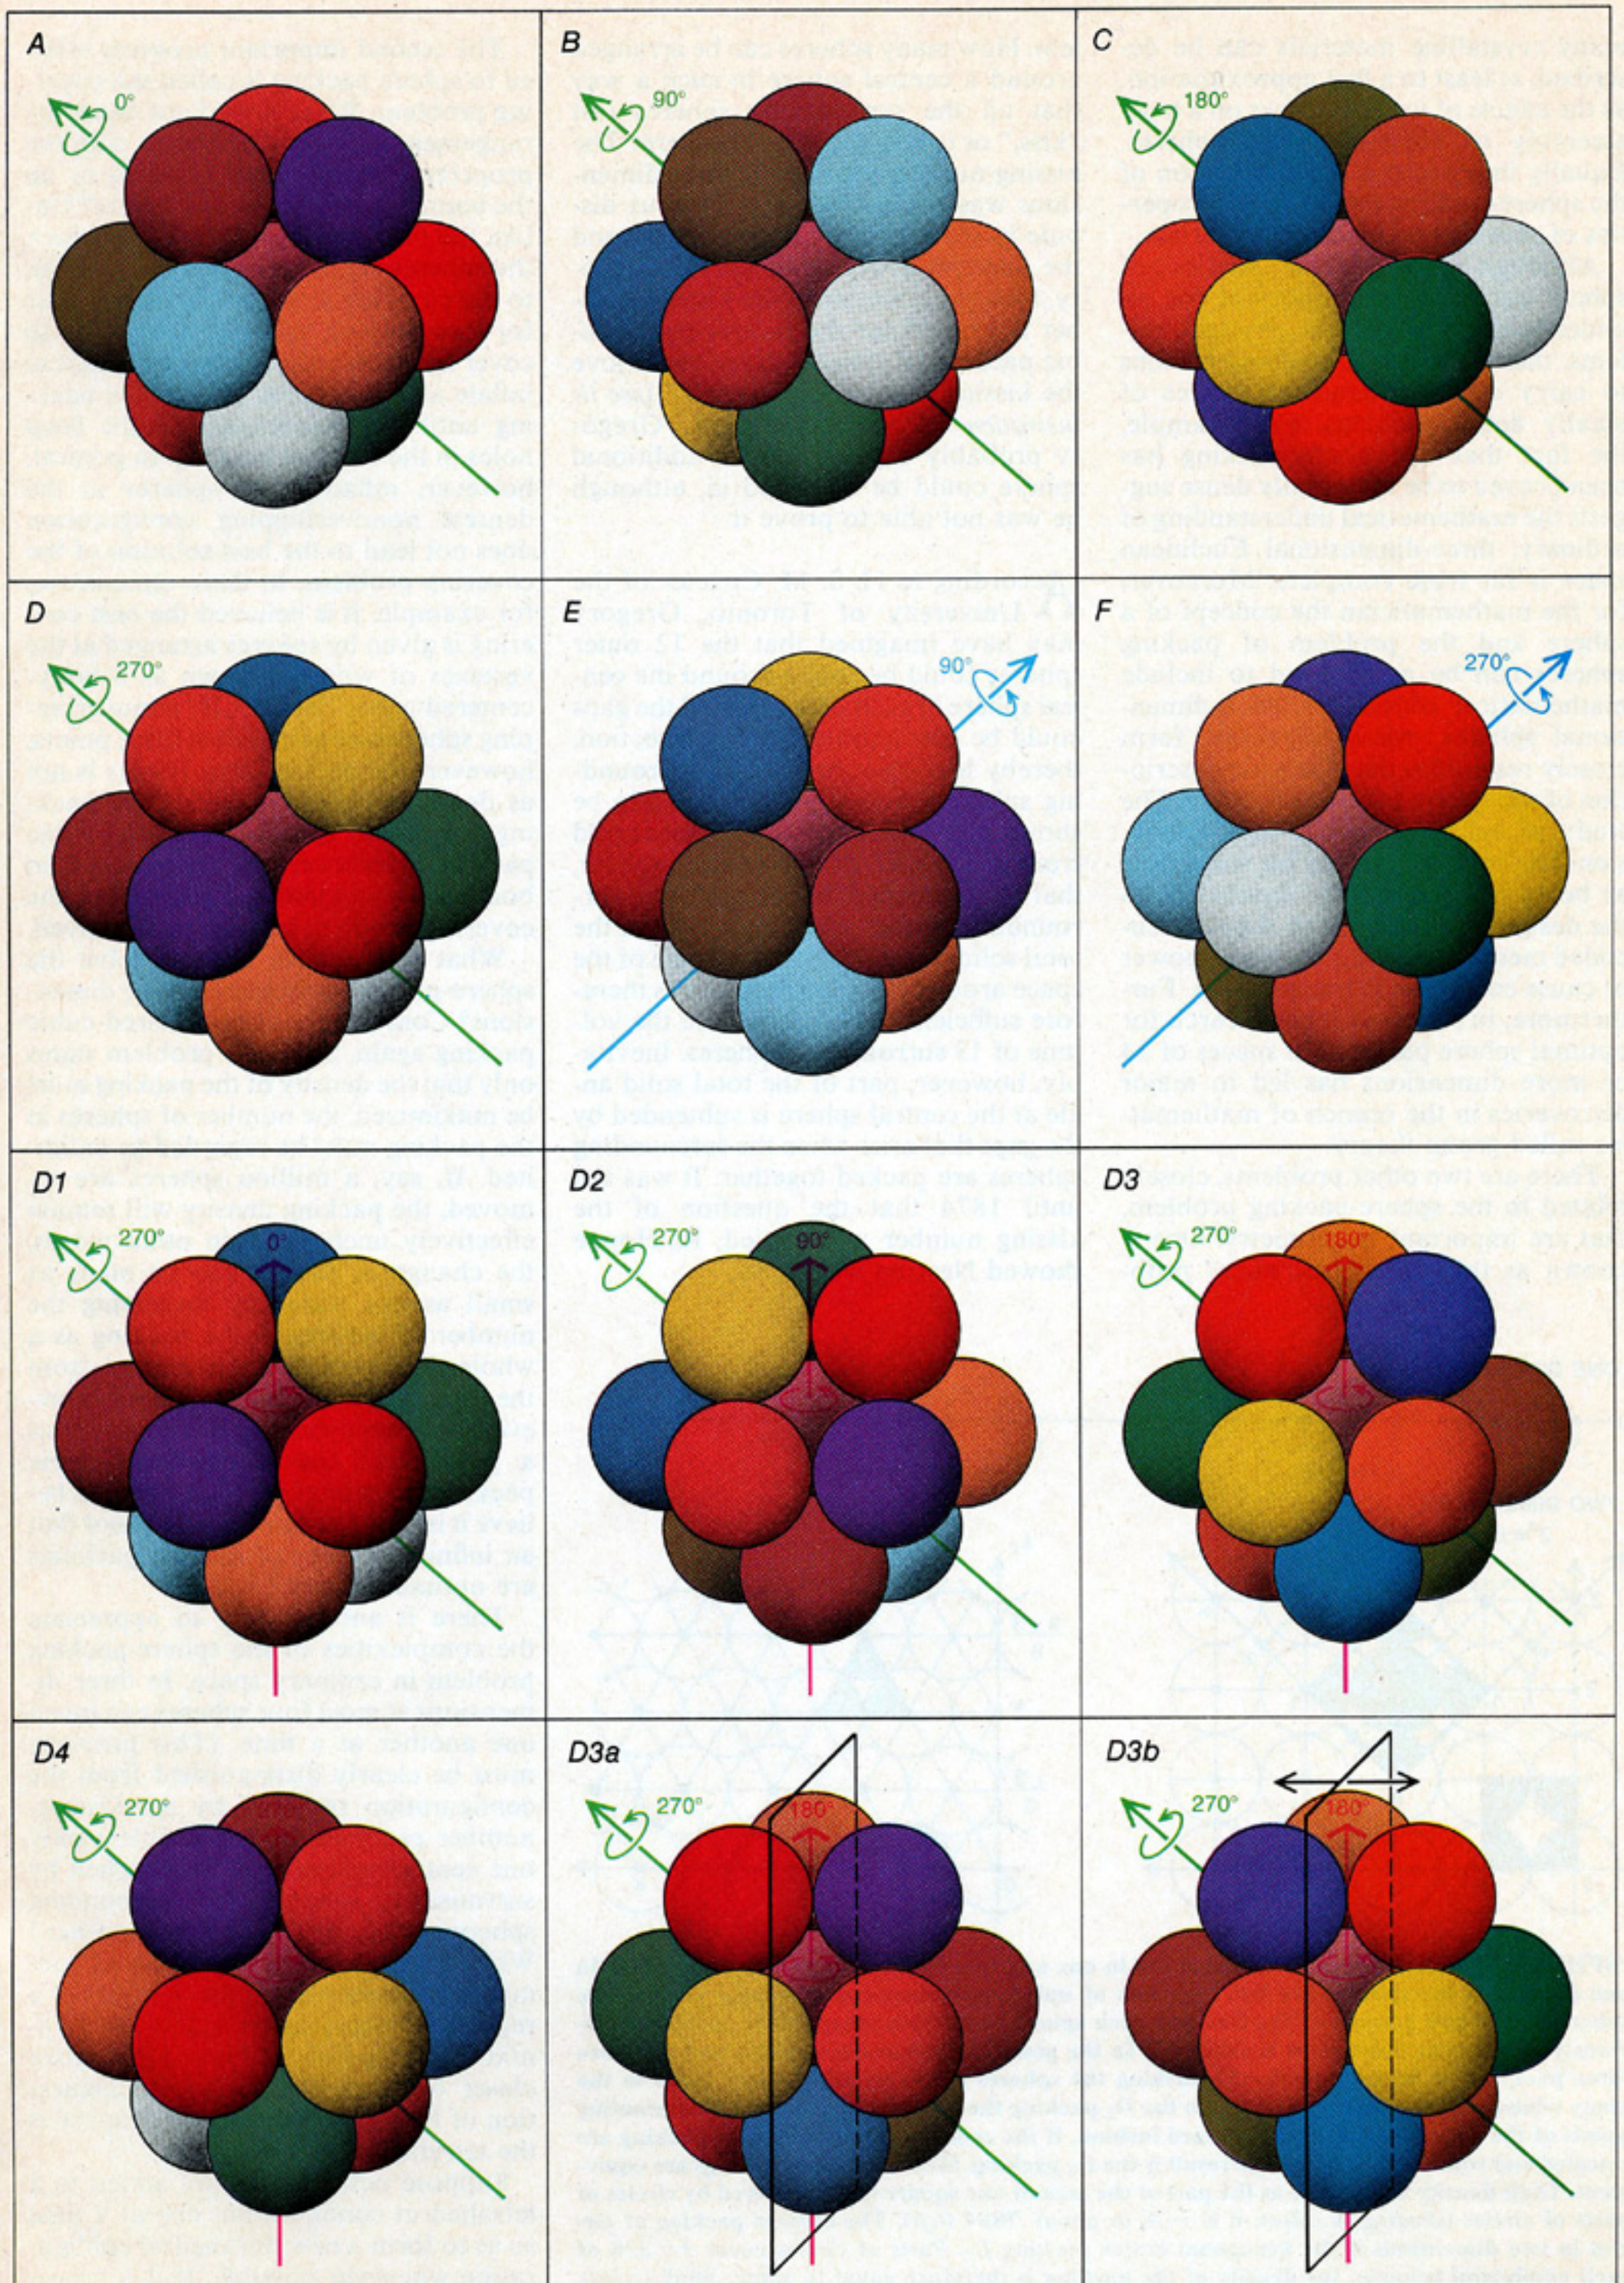
I am sorry to report, however, that this density has never been mathematically proved to be maximal. The least upper bound on the density obtained so far was found in 1958 by C. A. Rogers of the University of Birmingham, who proved that no packing of spheres can have a density greater than about .7796. The result is not particularly helpful to anyone looking for a more efficient way to pack ball bearings. Rogers' proof offers no construction of any sphere pack-

ing that comes close to his bound, and in the paper announcing the proof he remarked that "many mathematicians believe, and all physicists know," the correct answer is about 74 percent. In the quarter century since Rogers made the remark his assessment stands unchanged; the sphere-packing problem, so simple to state and so difficult to solve, remains one of the basic unresolved problems in mathematics.

Densely arranged configurations of hard spheres have been studied for many years in part because of their broad implications for an understanding of the behavior of solids and liquids. For example, the molecular properties of

FACE-CENTERED-CUBIC packing of spheres, often seen in fruit stands or in piles of cannonballs at war memorials, is thought to be the densest packing of spheres in three-dimensional space. In spite of centuries of effort, however, a proof of its maximal density has never been given. Each sphere in the packing "kisses," or touches, 12 other spheres; a proof that this number is maximal was not given until 1874. If the center of one sphere is fixed, the set of all possible rotations and reflections that permute the 12 surrounding spheres is called the symmetry group of the packing. The symmetry group of the face-centered-cubic packing has 48 elements; they can best be understood if the center of each sphere is thought of as the vertex of the polyhedron at the right, which is called a cuboctahedron. Any one of the six square faces of the cuboctahedron can become the front face by an appropriate rotation of the figure about the green or the blue axis (A-F). Each set of four spheres that form a square face (say the fourth face) can then assume one of four configurations if the entire figure is rotated about the red axis (D1-D4). Finally, each configuration (say the third one) can be reflected about a vertical plane to give a new configuration (D3a, D3b). The total number of elements in the symmetry group is therefore $6 \times 4 \times 2$, or 48. An analogous symmetry group, which describes the rotations and reflections of a dense packing of spheres in 24-dimensional space discovered by John Leech at the University of Glasgow, has been important in the mathematical theory of finite groups.





many crystalline materials can be described, at least to a first approximation, as the effects of various forces on a huge assembly of closely packed spheres. Equally important is the application of the sphere-packing model to the properties of powders and porous materials.

Although experimental studies of sphere packings are important for an understanding of certain physical systems, there are also compelling reasons to carry out mathematical studies of ideally dense packings. For example, the fact that no sphere packing has been proved to be maximally dense suggests the mathematical understanding of ordinary, three-dimensional Euclidean space is far from complete. Moreover, for the mathematician the concept of a sphere and the problem of packing spheres can be generalized to include mathematical objects called n -dimensional spheres, whose algebraic form closely resembles the algebraic description of the sphere in ordinary space. The study of sphere packings in n dimensions has been recognized for some time as being mathematically equivalent to the design of a finite set of digitally encoded messages that do not waste power or cause confusion in transmission. Furthermore, in recent years the search for optimal sphere packings in spaces of 24 or more dimensions has led to major discoveries in the branch of mathematics called group theory.

There are two other problems, closely related to the sphere-packing problem, that are important in geometry. One is known as the "kissing number" prob-

lem: How many spheres can be arranged around a central sphere in such a way that all the surrounding spheres just "kiss," or touch, the central sphere? The kissing-number problem in three dimensions was the subject of a famous dispute in 1694 between Isaac Newton and the Scottish astronomer David Gregory. Newton maintained the kissing number is 12, and for the face-centered-cubic packing of spheres described above the kissing number is indeed 12 [see illustration on preceding page]. Gregory probably argued that an additional sphere could be squeezed in, although he was not able to prove it.

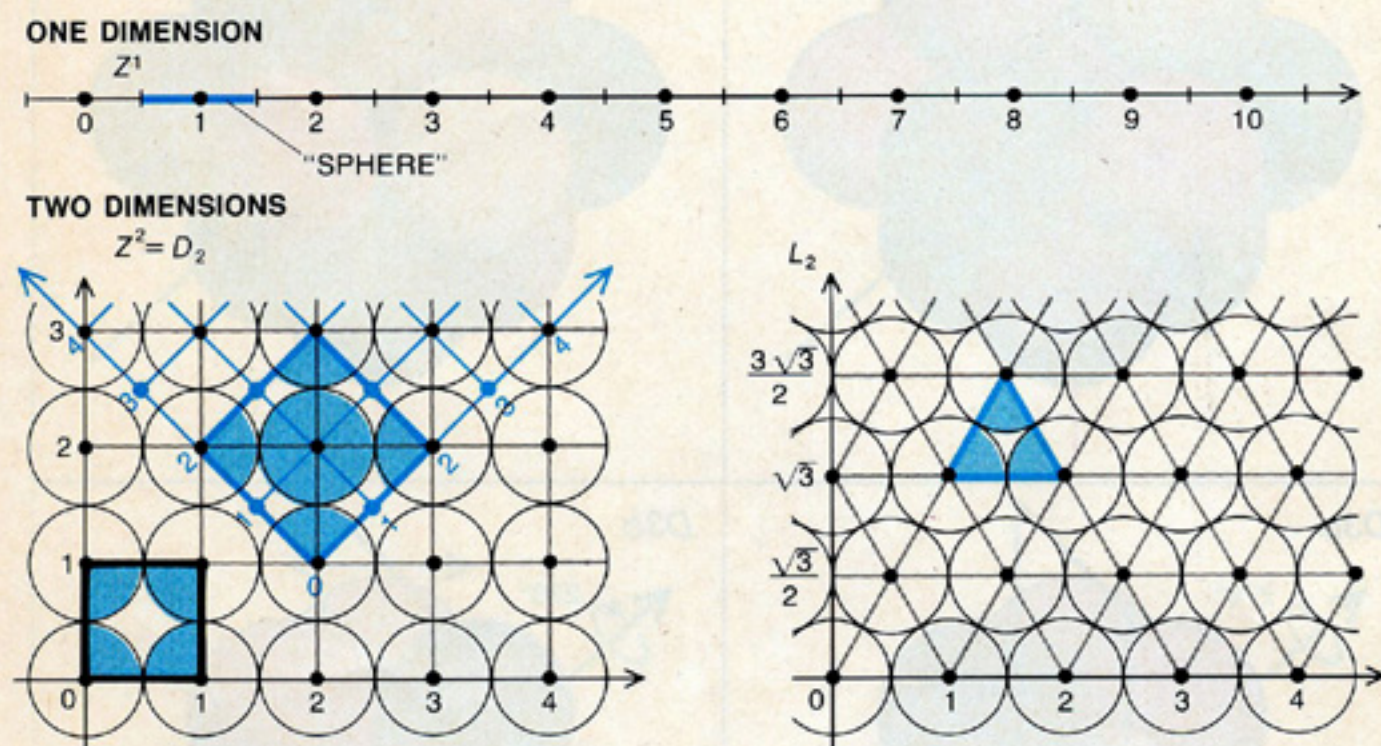
According to H. S. M. Coxeter of the University of Toronto, Gregory may have imagined that the 12 outer spheres could be rolled around the central sphere in such a way that all the gaps could be concentrated in one direction, thereby leaving room for 13 surrounding spheres. Actually it can readily be shown that the solid angle, measured from the center of the central sphere, that is subtended by one of the surrounding spheres is less than $1/13$ of the total solid angle. The total volume of the space around the central sphere is therefore sufficient to accommodate the volume of 13 surrounding spheres. Inevitably, however, part of the total solid angle at the central sphere is subtended by the gaps that arise when the surrounding spheres are packed together. It was not until 1874 that the question of the kissing number was settled; R. Hoppe showed Newton was right.

The second important problem related to sphere packing is called the covering problem: What is the least dense arrangement of identical spheres with the property that any point is inside or on the boundary of at least one sphere? Unlike the first packing problem, in which the spheres cannot overlap, the solution to the covering problem inevitably calls for overlapping spheres. One way to cover an entire volume with spheres is to inflate all the spheres in a sphere packing until they encompass all the deep holes in the original packing. In general, however, inflating the spheres in the densest nonoverlapping configuration does not lead to the best solution of the covering problem. In three dimensions, for example, it is believed the best covering is given by spheres arranged at the vertexes of what is known as a body-centered-cubic lattice. If nonoverlapping spheres are centered at these points, however, the packing that results is not as dense as other known sphere packings, such as the face-centered-cubic packing. Moreover, the conjecture that body-centered-cubic packings solve the covering problem has not been proved.

What makes it so hard to solve the sphere-packing problem in three dimensions? Consider the face-centered-cubic packing again. Since the problem states only that the density of the packing must be maximized, the number of spheres in the packing must be regarded as unlimited. If, say, a million spheres are removed, the packing density will remain effectively unchanged; in other words, the change in density can be made as small as one wants by increasing the number of spheres in the packing as a whole. For each sphere removed from the face-centered-cubic packing, however, a distinct packing is created. Thus a proof that the face-centered-cubic packing is optimal (as many people believe it is) is at the same time a proof that an infinite number of distinct packings are optimal as well.

There is another way to appreciate the complexities of the sphere-packing problem in ordinary space. In three dimensions at most four spheres can touch one another at a time. (This property must be clearly distinguished from the configuration required by the kissing-number problem. In that problem only one central sphere must be touched by surrounding spheres; the surrounding spheres need not touch one another.) When four spheres are mutually in contact, their centers form the vertexes of a regular tetrahedron, or triangular pyramid. Since the four spheres cannot move closer together, the densest configuration of four spheres possible in space is the tetrahedral configuration.


Suppose new spheres are added to a tetrahedral configuration one at a time so as to form a new tetrahedral configuration whenever possible. If this proce-




"SPHERE" PACKINGS can be carried out in one and two dimensions as well as in three. In one dimension the spheres are line segments of unit length, centered at integer points. The spheres cover 100 percent of the line and each sphere kisses two others. This packing is obviously as dense as possible; it is called Z^1 . In the plane the spheres are circles and there are three packings of interest. In the Z^2 packing the spheres are centered at every point in the plane whose coordinates are integers; in the D_2 packing the spheres are centered at alternating points of the Z^2 packing, in checkerboard fashion. If the coordinate axes of the Z^2 packing are rescaled and rotated 45 degrees, the result is the D_2 packing. Hence the two packings are equivalent. Their density is the same as the part of the area of one square that is covered by circles or parts of circles (shading in color); it is $\pi/4$, or about .7854 (left). The densest packing of circles in two dimensions is the hexagonal lattice packing L_2 . Parts of circles cover $\pi\sqrt{3}/6$ of each equilateral triangle; the density of the packing is therefore equal to about .9069 (right).

LAMINATED LATTICE PACKING L_3



BOTTOM LAYER



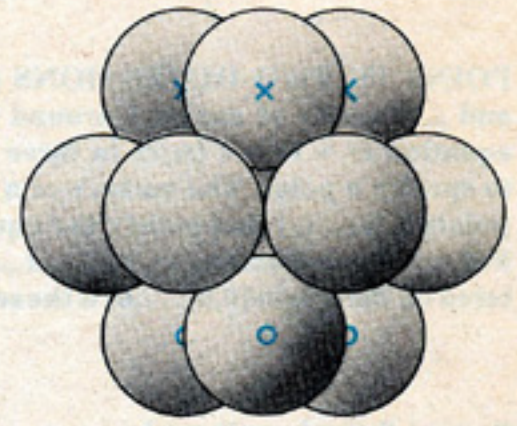

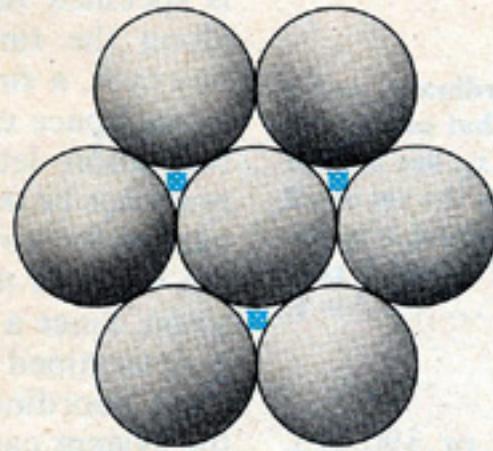

MIDDLE LAYER



TOP LAYER



HEXAGONAL CLOSE PACKING (NONLATTICE)



CONSTRUCTION of the laminated lattice packing L_3 , which is equivalent to the face-centered-cubic packing D_3 , is done by fitting together layers of spheres whose centers are arranged according to the hexagonal lattice packing L_2 . If the spheres are packed so that

the spheres in the third hexagonal layer are directly above the spheres in the first layer, the packing is called the hexagonal close packing. The hexagonal close packing is just as dense as L_3 , but the centers of the spheres in the packing do not satisfy the definition of a lattice.

ture could be continued indefinitely, the resulting arrangement of spheres would have the greatest packing density possible, since every tetrahedral configuration would be packed as densely as possible. Rogers' upper bound for the densest sphere packing possible is based on precisely this argument; a straightforward computation in spherical trigonometry shows that about .7796 of the volume of a tetrahedron is filled by the four nonoverlapping spheres that can be centered at its vertices. (The exact result is $\sqrt{2} [3 \arccos(1/3) - \pi]$.) Unfortunately, however, tetrahedrons do not fit perfectly together to fill space. The strategy of packing spheres into tetrahedral arrangements whenever possible, which might be called a greedy algorithm, eventually forces a bad move. At a certain stage in the procedure the growing configuration of spheres presents a surface that cannot accrete more spheres without wasting interior space. Hence although the greedy algorithm generates an optimal sphere packing over the short range of, say, a few spherical diameters, it turns out that the algorithm gives rise to a packing that is less dense on a global scale than the face-centered-cubic packing.

In order to make further progress on the three sphere-packing problems mathematicians have found it convenient to supplement geometric intuition with an analytic representation of the spheres in terms of their rectangular co-

ordinates. It is well known that any point in the plane can be specified by two coordinates, a horizontal coordinate x and a vertical coordinate y ; the point is generally written as the ordered pair (x,y) . For example, the point $(3,4)$ refers to the point in the plane three units to the right of the origin along the x axis and four units above the origin along the y axis.

The distance between, say, the point $(3,4)$ in the plane and any other point (x,y) can be calculated from the Pythagorean relation among the sides of a right triangle: The square of the distance between the two points is equal to the square of the distance between them along the x axis, $(x-3)^2$, plus the square of the distance between them along the y axis, $(y-4)^2$. Since a circle is by definition the set of all points in the plane equidistant from a central point (a,b) , any point (x,y) on the circumference must satisfy the equation $(x-a)^2 + (y-b)^2 = R^2$, where R is the radius of the circle. If the radius of the circle is equal to 1 and the center is at the origin $(0,0)$, the equation is much simpler: all points (x,y) on the circumference must satisfy the equation $x^2 + y^2 = 1$.

Similarly, any point in three-dimensional space is specified by three coordinates, x , y and z ; more suggestively, the point can be written (x_1, x_2, x_3) . The surface of a sphere of radius 1 centered at the origin is made up of all the points (x_1, x_2, x_3) such that $x_1^2 + x_2^2 + x_3^2 = 1$. The equation arises, much as it does


in two dimensions, from the geometric definition of a sphere and from two applications of the Pythagorean relation.

In more than three dimensions geometric intuition is of little value and one must begin to think exclusively in terms of coordinates. For example, a "point" in four-dimensional space is a mathematical object that requires four distinct numbers in order to be specified unambiguously; such a point is written (x_1, x_2, x_3, x_4) . If a person's height, weight, age and income are sufficient to unambiguously pick out the name of that person from a list, the four quantities can be regarded as specifying a point in a four-dimensional space.


A four-dimensional sphere is defined by analogy with the definitions for the circle and the sphere in two and three dimensions. All the points (x_1, x_2, x_3, x_4) on the "surface" of the sphere are found at some distance R from a central point (a_1, a_2, a_3, a_4) . The sum of the squares of the distance along each independent coordinate axis between any point (x_1, x_2, x_3, x_4) on the sphere and the central point (a_1, a_2, a_3, a_4) must always be equal to R^2 .

There has been a great deal of nonsense written in science fiction and elsewhere about the mysteries of the fourth dimension. In mathematical discussions one must not assume, as the physicist does, that the fourth dimension represents time. Furthermore, one must avoid the temptation to reify the somewhat metaphorical terms "surface,"

TWO DIMENSIONS



THREE DIMENSIONS



POINT IN TWO DIMENSIONS is specified by assigning a value to two coordinates, say x and y . A circle of radius 1 around the origin $(0,0)$ is the set of all points (x,y) that satisfy the equation $x^2 + y^2 = 1$ (left). In three dimensions three coordinates, say x_1, x_2 and x_3 , are needed to specify a point. The surface of a sphere of radius 1 around the origin $(0,0,0)$ is the set of all points (x_1, x_2, x_3) that satisfy the equation $x_1^2 + x_2^2 + x_3^2 = 1$. In n dimensions a point is specified by n coordinates, x_1, x_2, \dots, x_n . The surface of an n -dimensional sphere of radius 1 centered on the origin $(0,0, \dots, 0)$ is the set of points (x_1, x_2, \dots, x_n) such that $x_1^2 + x_2^2 + \dots + x_n^2 = 1$.

“point,” “sphere” and so on, which are applied to mathematical objects whose algebraic properties are otherwise perfectly straightforward. The terms are justified because the objects they refer to are constructed by analogy with the algebraic properties of ordinary circles and spheres; it must not be supposed, however, that the objects represent real geometric objects in some universe wider than our own. To repeat, in mathematics four-dimensional space consists of points with four coordinates instead of three (and the same holds for any number of dimensions).

Once the concept of a sphere is understood as a relation among coordinates, there are several problems, seemingly unrelated to sphere packing, that can be reduced to one of the three sphere-packing problems. The principal applications of high-dimensional sphere packing are to problems in digital communications, particularly in the construction of signals for use on a noisy channel and in the design of analogue-to-digital converters. It must be admitted, however, that most discussions of these applications have been theoretical; only quite recently have sphere packings been used to design practical systems. As communications systems become more sophisticated further applications can be expected.

In digital communications one of the major aims of the designer is to construct a list of distinct coded symbols, or code words, that can be transmitted with maximum reliability and minimum power. Each code word might be represented as, say, an eight-digit symbol, each digit of which can take on one of five distinct values: 0, 1/2, 1, -1/2 or -1. At first it would seem that the sys-

tem would provide for 5^8 , or 390,625, different code words, but the difference between many pairs of these code words is so small that such a system would be highly subject to random errors in transmission or to electrical interference. For example, the difference between the code word $(1,1,1,1,1,1,1,1)$ and the code word $(1,1,1,1,1,1,1,1/2)$ is far too small. If both code words were in use, it is likely they would often be confused. Another way to put the same point is that if the difference between two code words were as small as the difference between $(1,1,1,1,1,1,1,1)$ and $(1,1,1,1,1,1,1,1/2)$, a very large amount of power would be needed to guarantee that these two code words could be distinguished in the presence of background noise.

There is a general mathematical relation between the distinguishability of code words and the power needed to transmit them reliably. The relation was first formulated in 1948 by Claude E. Shannon, then of the Bell Telephone Laboratories, in his paper *A Mathematical Theory of Communication*. As my distinguished colleague David Slepian has said, “probably no single work in this century has more profoundly altered man’s understanding of communication” than Shannon’s paper. What Shannon showed is that given some fixed, finite amount of power there always exists a system of code words that can be transmitted essentially without error. The only qualification is that the rate at which the code words are transmitted cannot exceed a critical threshold called the capacity of the transmission channel. Unfortunately Shannon’s theorem is nonconstructive; it proves that such systems for encoding a signal exist, but it gives no hint of how they might be designed. Although many signaling sys-

tems have been constructed, schemes that perform as well as Shannon’s theorem promises have still not been found.

One way to design a signaling system that comes close to meeting the standards of Shannon’s theorem is to represent each signal as a point in n -dimensional space. For example, consider any sequence of eight numbers in the signaling system described above. Physically each of the numbers corresponds to a voltage level on a transmission line, and so each code word can be plotted on a two-dimensional graph as a series of eight distinct pulses whose height is specified for each of eight intervals along the time axis. Mathematically, however, a single point in eight-dimensional space can represent the same information: let the first number in each sequence be the value of the first coordinate of the point, the second number be the value of the second coordinate and so on. Since a point in eight dimensions is determined by fixing the values of all eight coordinates, every code word in the system can be represented as a distinct point in eight-dimensional space.

When the code words are represented as points, two important features of a system of code words can be given a geometric interpretation. First, remember that the code words must be reliably distinguishable from one another. In eight-dimensional space this suggests that the points representing the possible code words must be separated by a certain minimum “distance.” How well does the Pythagorean distance between two points in eight dimensions measure the distinguishability between two code words? In order to determine the Pythagorean distance the difference between the two values of each coordinate of the two points must be squared. Small differences between coordinate values (that is, differences less than 1) are thereby reduced, whereas large differences (that is, differences greater than 1) are magnified. Because a small difference in the voltage levels between two signals is much more likely to lead to confusion between code words than a large difference, the Pythagorean distance is a reasonable measure of distinguishability.

For example, consider the two code words $(1,1,1,1,1,1,1,1)$ and $(1/2,1/2,1/2,1/2,1/2,1/2,1/2,1/2)$. The square of the distance between the two points is the sum of eight squares, each of the form $(1 - 1/2)^2$, and so the distance between them is $\sqrt{2}$. According to the Pythagorean measure of distance, the two code words $(1,1,1,1,1,1,1,1)$ and $(0,0,1,1,1,1,1,1)$, which differ only in the first two coordinates, are also $\sqrt{2}$ units apart, and so they are just as easy to distinguish as $(1,1,1,1,1,1,1,1)$ and $(1/2,1/2,1/2,1/2,1/2,1/2,1/2,1/2)$.

A second important feature of any system of code words is that the power

needed to transmit them must be minimized. It follows from the elementary definitions of power and voltage that power is directly proportional to the square of voltage; for a simple circuit the power is the square of the voltage divided by the resistance in the circuit. The total power necessary to transmit an eight-digit code word is consequently the sum of the squares of all eight digits that make up the code word. This sum is the square of the distance between the point representing the code word in eight-dimensional space and the origin $(0,0,0,0,0,0,0,0)$.


The design of a signaling system that is reliable and makes efficient use of power can therefore be reduced to the geometric problem of placing points inside a region of space while constraining them not to be too close together. If the points must be at least a distance of, say, $\sqrt{2}$ apart, the problem is equivalent to the problem of finding the densest packing of spheres whose radius is half that distance, $\sqrt{2}/2$. A closely related problem is to find a set of code words that all have the same energy. This problem is equivalent to the problem of placing as many points as possible on the surface of an n -dimensional sphere while constraining them not to be too close together. That problem is in turn a generalized version of the kissing-number problem.

It turns out that in a space of eight dimensions there is an extremely dense packing called the E_8 packing; it was discovered in the last third of the 19th cen-

DESIGN OF A CODE for the efficient transmission of information is closely related to the sphere-packing problem. The code is to be a finite set of signals called code words that are easily distinguished from one another and do not waste electric power. If each code word is a sequence of, say, three discrete voltage levels, each sequence can be represented as a point in three-dimensional space: the first coordinate of the point is the numerical value of the first voltage level, the second coordinate is the value of the second voltage level and so on (A-D). The transmission power required for each voltage pulse is proportional to the square of the voltage, and so the total power needed to transmit one code word is the sum of the squares of the three discrete voltages associated with the code word. The sum is equal to the square of the distance from the origin to the point in three-dimensional space that represents the code word. Thus the problem of minimizing transmission power is equivalent to the problem of placing all the points that represent code words as close to the origin as possible. On the other hand, the need to distinguish code words from one another can be treated as the requirement that the points in space representing the code words not be closer together than some minimum distance d . Meeting the two requirements simultaneously is geometrically equivalent to packing hard, nonoverlapping spheres of radius $d/2$ around the origin as densely as possible (E).

tury by the Russian mathematicians Alexander N. Korkin and E. I. Zolotareff and by the English lawyer and amateur mathematician Thorold Gosset. The centers of the spheres in the E_8 packing are all the points whose coordinates are equal to whole numbers or to whole numbers plus a half. The sum of the

coordinates for each point must be an even number. There are 240 such points whose distance is $\sqrt{2}$ from the origin: 112 points of the form $(\pm 1, \pm 1, 0, 0, 0, 0, 0, 0)$, where the two 1's and any combination of signs can appear in any positions, and 128 points of the form $(\pm 1/2, \pm 1/2, \pm 1/2, \pm 1/2,$



$\pm 1/2, \pm 1/2, \pm 1/2, \pm 1/2$), where the number of minus signs is even.

The E_8 packing could become the basis of a practical and efficient signaling scheme. If the scheme were to include exactly 240 code words, the 240 points in the E_8 packing that are equidistant from the origin could be chosen as the code words. In most practical systems, however, the number of distinct code words needed is some integral power of 2. For example, in medium-distance telephony the digital transmission system known as pulse-code modulation is now in wide service. The voltage of the voice signal is measured every $1/8,000$ second, and the measured samples are quantized, or replaced, by one of 2^8 , or 256, levels. Each quantized voltage level is then expressed as an eight-digit binary number; such binary numbers make up what is called the source code.

It has been known since the work of Harry Nyquist of Bell Laboratories in the 1920's that the voice signal can be reconstructed from the sample values alone. The voice signal can be thought of as a quantity, such as air pressure or voltage, that varies continuously with time. Early in the 19th century Jean Baptiste Joseph Fourier proved that the graph of any such quantity can be approximated to any desired degree of accuracy by superposing sine and cosine curves of appropriate amplitude and frequency. The sine and cosine curves (which, for an arbitrary curve,

may be infinite in number) are called the Fourier components of the curve on the graph.

Suppose the graph of a quantity can be exactly generated by superposing a finite number of Fourier components whose frequencies are no higher than some frequency W cycles per second. Then what Nyquist showed is that the graph can also be perfectly reconstructed solely from the values it takes every $1/2W$ second. For example, a voice signal that has no frequency components higher than 4,000 cycles per second can be perfectly reconstructed from samples of the signal taken every $1/8,000$ second. It is therefore enough to sample the voice signal and transmit only the sample values, which are represented by the code words in the source code, instead of transmitting the entire voice signal. If the sample values were transmitted without being rounded off, Nyquist's theorem shows that an entire second of speech signal, which corresponds to 8,000 sample values, can be represented as a single point in a space of 8,000 dimensions. This shows the power of mathematics.


For efficient transmission the source-code numbers that represent the sample values must be further encoded by another code, called a channel code; it is the channel code that pertains to the sphere-packing problem. An excellent way to derive a channel code from the

E_8 sphere packing is to encode each pair of successive eight-digit binary numbers in the source code. Each 16-digit binary number that results is then assigned to the center point of one of 2^{16} , or 65,536, spheres in the E_8 packing. A good channel code can then be constructed by choosing the 65,536 center points that lie closest to the origin. At the receiving end of the telephone line the code words that correspond to the coordinates of each center point are converted back into the binary numbers of the source code, and the voice signal is reconstructed from the binary numbers.

There is a second major application of sphere packing in digital communications that I shall describe briefly. Remember that in deriving the binary numbers in the source code from the voice signal of the telephone it is necessary to quantize the precise intensity of the signal to one of 256 levels. The real world is full of awkward numbers such as .7913..., but the world of computers and digital systems must ultimately deal only with round numbers such as 0 and 1. Any device that rounds off a continuously variable quantity to some set of discrete values is called an analogue-to-digital converter, or quantizer.

Quantization can be carried out in two or more dimensions as well as along a single coordinate axis. Imagine that the plane is divided into regions, not necessarily congruent, and imagine that within each region one point has been marked. Any such array of points and regions can function as a two-dimensional quantizer; the input to the quantizer is a pair of real numbers that specify some arbitrary point, and the output is the preselected quantizing point that lies in the same region of the plane as the arbitrary point. Thus any point in the plane is rounded off to one of the quantizing points. The process compresses input data; a single indexing number for the quantizing point can be transmitted in place of the precise coordinate values of the data point.

Quantizing introduces errors, and so one tries to choose the quantizing points in such a way as to minimize the average error. For example, if the input to the quantizer is uniformly distributed, or in other words if every input value is equally probable, it is straightforward to calculate the average error for a number of quantizing schemes. If a single coordinate axis is divided into equal segments of unit length and the quantizing points are centered on each segment, the average error is $1/12$, or about .0833. The same data can be quantized in two dimensions: the data points are paired, and each pair is considered a point in the plane. If the plane is then divided into squares and the quantizing points are centered on each square, the average error is still $1/12$. On the other hand, if the plane is divided into regular



QUANTIZING DATA from a continuously variable source is closely related to the problem of covering space with the least dense arrangement of overlapping spheres. For example, in order to quantize data in two dimensions the input data are paired and each pair is treated as the coordinates of a point in the plane. Each data point, such as the point *A*, is then rounded off to a quantizing point, such as the point *B*, that lies in the same preselected region of the plane as the data point. The problem is to choose the quantizing points and the partition of the plane in a way that minimizes the average quantizing error. If the data points are uniformly distributed and the quantizing points are chosen as the centers of squares, the average error is $1/12$. A better way of quantizing can be derived from the best covering of the plane by circles. That covering is generated from the hexagonal packing of circles (black circles) by increasing the radius of each circle by just enough to include every point in the plane inside or on the circumference of at least one circle (colored circles). The radius *R* of the covering circles is the distance from the center of a circle to the nearest "deep hole" in the packing. If the deep holes are connected by the appropriate straight lines, the plane is partitioned into regular hexagons; for uniformly distributed data points the average quantizing error for quantizing points that are centered on regular hexagons is $5\sqrt{3}/108$, or about .0802, which is slightly less than $1/12$.

hexagons having the same area as the squares, and if each quantizing point is at the center of a hexagon, the average quantization error can be reduced to $5\sqrt{3}/108$, or about .0802.

Remarkably, it turns out that a corresponding improvement can always be made even if the data points are not evenly distributed. In a 1963 doctoral dissertation at Stanford University, P. L. Zador showed that it is always possible to reduce the average error by quantizing in a space of higher dimensions. It is more efficient to wait until several data points have been collected and then quantize them all at the same time by regarding them as a point in n -dimensional space than it is to quantize them one at a time along a single axis. In quantizing it pays to procrastinate.

Unfortunately Zador's result, like Shannon's theorem, is nonconstructive. The problem of finding good multidimensional quantizers even for uniformly distributed data is still unsolved. There are several sphere packings, however, that appear to give rise to excellent quantizing schemes. Consider the packing of spheres in two dimensions, that is, the packing of circles. It has been known since 1940 that circles can be packed at maximum density if they are arranged so that each circle is surrounded by six others [see illustration on opposite page].

Imagine now that each circle in the packing is bounded by a thin, flexible membrane and that the interior of the circle is inflated. As the circles get bigger the membranes press against one another to fill the remaining space in the plane; if the inflation is uniform throughout the plane, each circle will expand to form a regular hexagon. As I have mentioned, basing the quantization of uniformly distributed data on hexagonal regions leads to the minimum average error. A similar expansion of the eight-dimensional spheres in the E_8 packing also leads to low quantization error, which is even smaller than the error in two dimensions. The general problem of quantization, which requires that space be broken up, or covered, by a discrete partition, is closely related to the problem of finding the best covering of space by spheres.

The search for dense sphere packings in multidimensional spaces is greatly simplified by focusing attention on certain kinds of packings called lattice packings that have a highly regular configuration. Consider the hexagonal packing of circles I have just described: notice that the centers of any two adjacent surrounding circles and the center of the circle in the middle form an equilateral triangle. In order to calculate the density of the packing it is sufficient to determine the proportion of each triangle that is covered by circles or parts of circles. Because the triangles tessellate,

or completely fill, the plane and because the configuration of circles within a triangle is always the same, the circle density in one triangle is equal to the density of the packing throughout the plane. By applying elementary geometry one can show the density is equal to $\pi\sqrt{3}/6$, or approximately .9069.


The foregoing calculation would not have been possible if no repeating unit that fills the plane could have been found. It is nonetheless easy to imagine that there exist highly irregular sphere packings that have no repeating units. Such packings are much harder to study than those that repeat: not only is the density of an irregular packing difficult or impossible to determine but also even the coordinates of all the center points may not be specifiable. The definition of a lattice packing guarantees that these disadvantages can be avoided. A sphere packing is said to be a lattice packing if whenever there are two spheres, one that has its center at the point (u_1, u_2, \dots, u_n) and the other at the point (v_1, v_2, \dots, v_n) , there are also spheres in the packing with centers at all points of the form $(au_1 + bv_1, au_2 + bv_2, \dots, au_n + bv_n)$, where a and b are any whole numbers. The center coordinates of the latter spheres are said to be generated from the center coordinates of the first two spheres.

The simplest lattice packing is the cubic lattice packing, in which the center coordinates of each sphere in the packing are integers; the cubic lattice in an arbitrary number of dimensions n is designated Z^n . The one-dimensional "cubic" lattice Z^1 is made up of discrete line segments, each of unit length, that are centered at the integer points on the line. The "spheres," or line segments, cover 100 percent of the line, and each sphere touches two others; hence Z^1 solves the sphere-packing problem and the kissing-number problem in one dimension.

In two dimensions, however, the square lattice Z^2 is not the densest packing. Its packing density is $\pi/4$, or about .7854, far less than the density of the hexagonal packing of circles [see illustration on page 118]. In a like manner the sphere-packing density of the cubic lattice Z^3 is relatively low: it is $\pi/6$, or about .5236. A much denser family of lattice packings can be generated from cubic lattices if the spheres are centered at alternate points of the lattice in checkerboard fashion. To construct the new family of lattices color the points of a cubic lattice alternately red and black and make the centers of the spheres coincide with the black points. Equivalently, the centers of the new packing are the points with integer coordinates that add up to an even number. In an arbitrary number of dimensions n this packing is designated D_n . In D_3 , for example, the origin $(0,0,0)$ and the

point $(1,1,0)$ are legal centers of spheres, but the point $(1,0,0)$ is not because $1 + 0 + 0$ is an odd number.

The sequence of lattice packings D_3 , D_4 , D_5 and so on is of considerable importance for the sphere-packing problem. The packing D_3 is the face-centered-cubic lattice. A model constructed of Ping-Pong balls shows that one repeating unit cell of the lattice is a cube two units on a side having a sphere at its center; the radius of each sphere is $\sqrt{2}/2$. The density of the packing can be cal-



LAMINATED lattice packings of spheres in n dimensions are built up by packing together layers of a suitable laminated lattice packing in the next smaller dimension $n - 1$. For example, the two-dimensional hexagonal packing L_2 can be built up by stacking rows of circles whose centers are arranged according to the Z^1 packing. Similarly, layers of spheres, each layer arranged according to the hexagonal lattice packing, can be stacked to yield the lattice packing L_3 , the densest packing known in three dimensions. John Horton Conway of the University of Cambridge and the author have continued the construction and found all the laminated lattices in dimensions up to 25. The laminated lattices in dimensions 1 to 10 and in dimensions 14 to 24 are unique; there are two laminated lattices in dimension 11, three each in dimensions 12 and 13, 23 in dimension 25 and at least 75,000 in dimension 26.

culated by finding the fraction of the volume of the cube that is filled by spheres; it is equal to $\sqrt{2}\pi/6$, or about .7405. Although it is possible that denser sphere packings exist in three dimensions, Carl Friedrich Gauss proved in 1831 that D_3 is the densest three-dimensional lattice packing. It is also known that D_4 and D_5 are the densest lattice packings in four and five dimensions.

Above five dimensions, however, D_n is not the densest lattice packing, and by the time one reaches D_8 there are huge gaps between the spheres. The gaps are so large that it is possible to slide another copy of D_8 into the holes without overlapping the spheres; the result turns out to be the E_8 lattice. In 1934 H. F. Blichfeldt of Stanford University proved that E_8 is the densest lattice packing in eight dimensions, and he showed that certain cross sections of E_8 called E_6 and E_7 are the densest lattice packings in six and seven dimensions. In these dimensions no denser packings, which would necessarily be nonlattice packings, have been found since.

In 1965 John Leech, then at the University of Glasgow, constructed a remarkable sphere packing in 24-dimensional space; his construction is outlined in the illustration below. The study of the Leech lattice, as the packing is usually called, has led to a deeper understanding of the properties of other higher-dimensional lattices and to important results in the theory of groups. The lattice is almost certainly the densest sphere packing possible in 24 dimensions. C. A. Rogers, arguing as he

did for sphere packing in three dimensions, gave bounds for the maximum density of packings in any n -dimensional space; his bound for any 24-dimensional sphere packing is only slightly greater than the density of the Leech lattice. Each sphere in the lattice touches 196,560 others, and in 1979 A. M. Odlyzko of Bell Laboratories and I proved that this number solves the kissing-number problem in 24 dimensions. The same method of proof also solved the kissing-number problem in eight dimensions; there the answer is 240, the number of spheres that touch one sphere in the E_8 lattice. These two results were found independently by V. I. Levenshtein of the L. V. Keldysh Institute of Applied Mathematics in Moscow. Incidentally, the problem remains unsolved in all other dimensions except one, two and three (where the answers are respectively two, six and 12).

The Leech lattice has been indispensable to group theorists for the construction of certain finite simple groups. These groups are the building blocks of all groups having a finite number of elements; their classification, which has just been completed, has preoccupied many mathematicians for more than 50 years. The simple groups play much the same role in group theory that the prime numbers play in number theory and the chemical elements play in chemistry. Several important simple groups have now been constructed by considering the set of all the rigid rotations and reflections of the Leech lattice that leave the central sphere fixed and permute the

surrounding spheres. This set of operations is called the symmetry group of the packing; the analogous symmetry group of the face-centered-cubic packing in three dimensions is shown in the illustration on pages 116 and 117.

The symmetry group of the Leech lattice was found in 1968 by John Horton Conway of the University of Cambridge. Its order, or the number of elements in the group, is immense, although not particularly large in the context of group theory: it is $2^{22} \times 3^9 \times 5^4 \times 7^2 \times 11 \times 13 \times 23$, or 8,315,553,613,086,720,000. From this group, which is not a simple group, Conway constructed three previously unknown simple groups whose order in each case exactly divides the order of the symmetry group of the Leech lattice. In 1981 the Leech lattice enabled Robert L. Griess, Jr., of the University of Michigan to construct one of the last finite simple groups to be found. It is rather larger than Conway's groups and has been nicknamed the monster: the number of elements in the group is $2^{46} \times 3^{20} \times 5^9 \times 7^6 \times 11^2 \times 13^3 \times 17 \times 19 \times 23 \times 29 \times 31 \times 41 \times 47 \times 71$, or 808,017,424,794,512,875,886,459,904,961,710,757,005,754,368,000,000,000. Griess's construction is not at all straightforward, however, and one of the fascinations of the Leech lattice is that one feels there should be a more direct connection between it and the monster simple group.

The Leech lattice is such a dense packing that its influence is felt in all lower dimensions. It is not surprising that a slice of a good packing gives a good packing in a space of one lower dimension; for example, one slice through D_3 exposes a surface of hexagonally packed spheres. Suitable cross sections of the Leech lattice, however, give rise to the densest known packings in all dimensions less than 24 except 10, 11 and 13. For example, one eight-dimensional slice through the Leech lattice exposes the E_8 lattice.

Since dense packings can be built from the Leech lattice, so to speak, from the top down, it is tempting to ask how the Leech lattice might be built from the bottom up, that is, from dense packings in lower dimensions. It turns out its construction can be done in a particularly simple way. Begin with the densest packing possible in one dimension, Z^1 . At the center of each one-dimensional sphere in Z^1 construct a two-dimensional sphere whose radius is $1/2$. Now construct another layer of two-dimensional spheres identical with the first layer and fit it into the holes in the first layer as tightly as possible. If an infinite number of layers are packed together in this way, the resulting two-dimensional lattice is the dense, hexagonal packing; because of this construction it can be

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	1	1	1	0	1	1
0	1	0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	0	0	1	1	1	0	1
0	0	1	0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	0	0	1	1	1	1
0	0	0	1	0	0	0	0	0	0	0	0	1	0	1	1	0	1	0	0	0	1	1	1
0	0	0	0	1	0	0	0	0	0	0	0	1	1	0	1	1	0	1	0	0	0	1	1
0	0	0	0	0	1	0	0	0	0	0	0	0	1	1	1	0	1	1	0	1	0	0	1
0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	1	0	1	1	0	1	0	0	1
0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	1	0	1	1	0	1	0	1
0	0	0	0	0	0	0	0	1	0	0	0	1	0	0	0	1	1	1	0	1	1	0	1
0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	1	1	1	0	1	1	1
0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0

CONSTRUCTION OF THE LEECH LATTICE, the densest known packing of spheres in 24-dimensional space, is based on the 24-digit binary sequences shown. The set of all possible sums of the 12 binary sequences, where the addition is carried out modulo 2, is made up of 2^{12} , or 4,096, binary sequences called code words. (In addition modulo 2 the sum of 1 and 1 is equal to 0; the digit ordinarily carried in binary arithmetic is ignored.) The 2^{12} code words make up an efficient code for transmitting information that was devised by Marcel J. E. Golay at the U.S. Army Signal Corps Engineering Laboratories in 1949. The centers of the spheres in the Leech lattice all have the form $2C + 4X$ or $I + 2C + 4Y$, where C is a code word of the Golay code, I is the point $(1,1,\dots,1)$ in 24 dimensions and X and Y range over all points in 24 dimensions whose coordinates are all integers. The sum of the coordinates of each point X must be even and the sum of the coordinates of each point Y must be odd. Each sphere has radius $2\sqrt{2}$, and the spheres closest to the origin have centers such as $(\pm 4, \pm 4, 0, \dots, 0)$, $(\pm 2, \pm 2, \pm 2, \pm 2, \pm 2, \pm 2, \pm 2, \pm 2, 0, \dots, 0)$ and $(\mp 3, \pm 1, \pm 1, \dots, \pm 1)$. Each sphere touches 196,560 others.

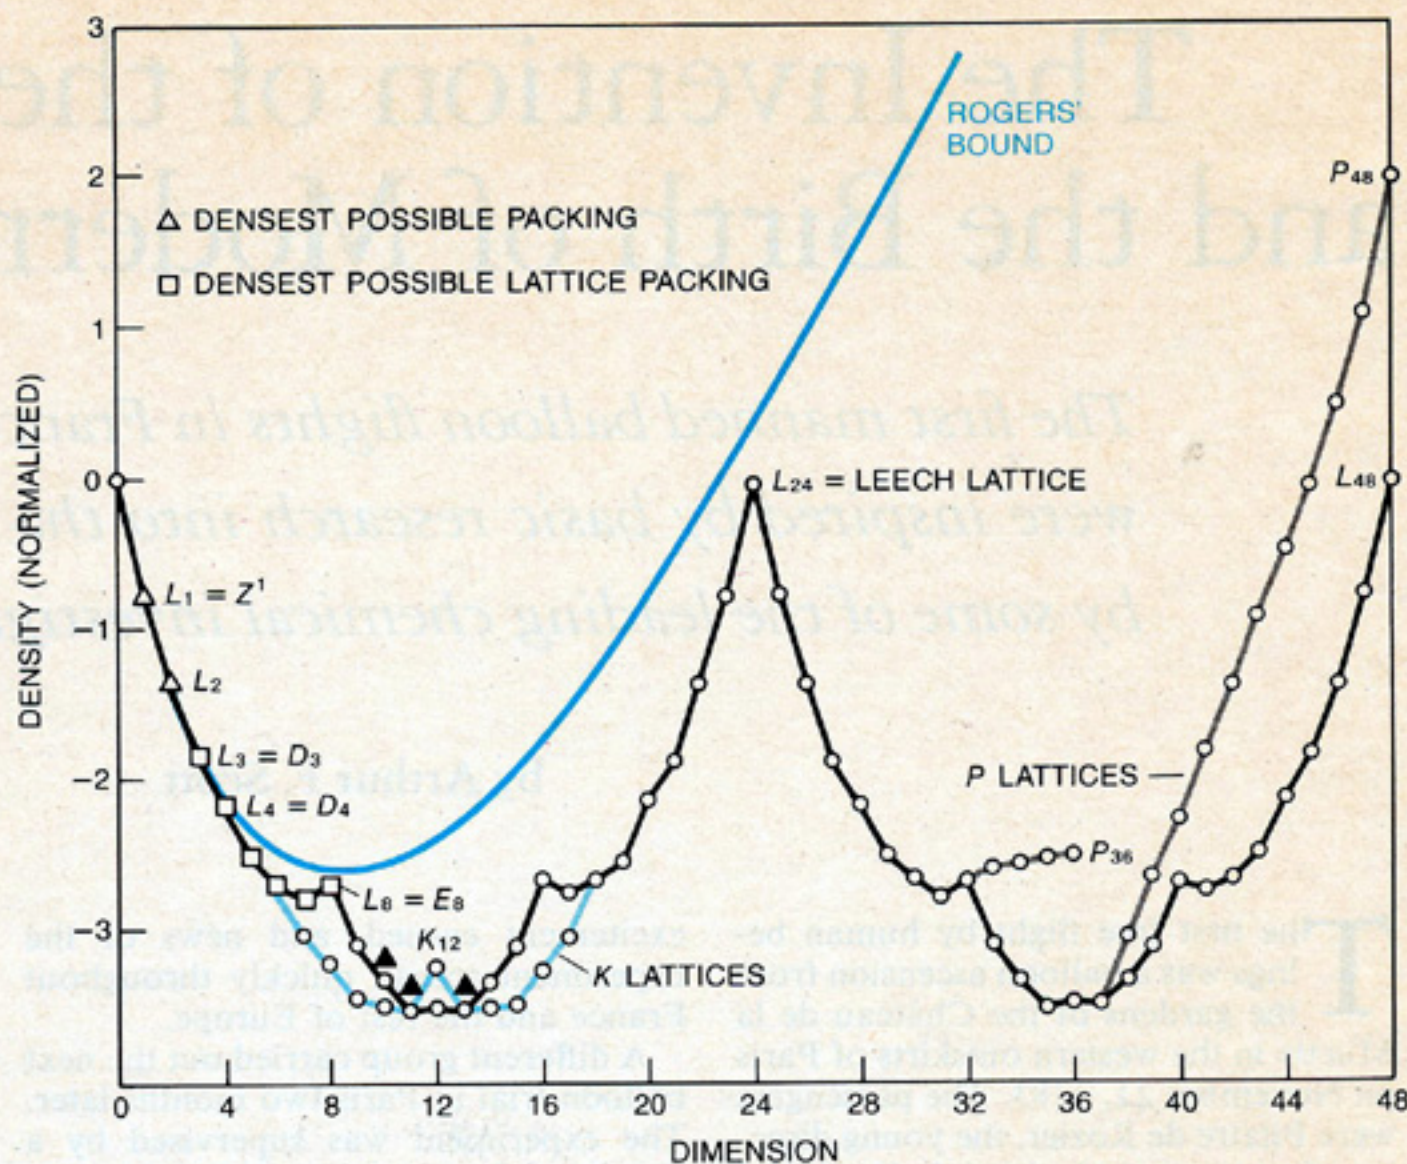
called the laminated packing in two dimensions, abbreviated L_2 .

I have already suggested how one goes to three dimensions: a ball of radius $1/2$ is emplaced at the center of each circle in L_2 , and identical layers of balls are fitted into the holes of the first layer in such a way as to form a lattice. Since this packing is equivalent to D_3 , the laminating procedure gives the best-known packing in three dimensions as well as in two. If the procedure is continued in higher dimensions, one dimension at a time, the resulting lattice packings are extremely dense. It has been known for a long time that L_4 and L_5 are equivalent to D_4 and D_5 , and that L_6 , L_7 and L_8 are equivalent respectively to E_6 , E_7 and E_8 . Hence the laminated lattices are the densest lattice packings possible in up to eight dimensions.

Recently Conway and I continued the laminating process a bit further: we found all the laminated lattices in up to 25 dimensions and examples of them for each dimension up to 48. It turns out that although there is a unique laminated lattice in every dimension up to 10, there are two distinct ways to stack layers of the L_{10} lattice. The two 11-dimensional lattices have the same density but different kissing numbers. There are three 12-dimensional and three 13-dimensional laminated lattices but only one such lattice for each dimension from 14 through 24. The laminated lattice L_{24} is the Leech lattice. It is only in spaces having 11, 12 and 13 dimensions where cross sections of the Leech lattice are known that have a higher density than the laminated lattices in those dimensions.

The laminated lattices are built, layer on layer, by fitting each new layer as snugly as possible into the deep holes of the preceding one. Hence the investigation of laminated lattices is closely related to the covering problem: enlarge every sphere in any packing just enough to encompass the deep holes, and by definition the packing must cover the space. In 1966 Leech conjectured that if each of the spheres in the Leech lattice is enlarged by a factor of $\sqrt{2}$, the spheres would cover all the points in the space; such a covering might well be the best-possible covering of 24-dimensional space. Leech's conjecture, however, was quite difficult to prove, primarily because of the complexity of the deep holes in the Leech lattice.

In the two-dimensional lattice L_2 it is clear there is only one kind of deep hole; the maximum distance from any point in the lattice is always a point bounded by three circles, and so there is nothing to distinguish one deep hole from any other. In the Leech lattice, however, Conway, Richard A. Parker and I classified 23 distinct kinds of deep hole before we were able to prove Leech's conjecture. Not surprisingly there turn out to



DENSEST KNOWN PACKINGS of spheres in spaces of up to 48 dimensions are plotted according to a method suggested by John Leech; the "normalized" density of the packing depends on the dimension in which the packing is done. Its definition is based on the fact that the density of the 24-dimensional Leech lattice, divided by the volume of a 24-dimensional sphere of unit radius, is equal to 1. (The volume of an n -dimensional sphere of radius 1 is equal to $\pi^{n/2}/(1 \times 2 \times 3 \times \dots \times n/2)$ if n is even, or to $2(2\pi)^{(n-1)/2}/(1 \times 3 \times 5 \times \dots \times n)$ if n is odd.) That quotient in an arbitrary space of dimension n is called the center density D . The normalized density shown on the graph is equal to $\log_2 D + n(24 - n)/96$; for laminated lattice packings the graph is symmetrical about the normalized density of the Leech lattice. The graph shows that L_3 , L_8 and L_{24} are quite close to the least upper bound known for the density of any sphere packing. The laminated lattice packings are the densest known packings in all dimensions up to 32 except dimensions 10 to 13. There is an alternative sequence of lattices called the K_n sequence, which starts at L_6 and rejoins the laminated sequence at L_{18} . The K_n packings are denser than the L_n packings in dimensions 11, 12 and 13. The P_n packings are also lattice packings, but the densest known packings in 10, 11 and 13 dimensions are nonlattice packings; all these packings are constructed from codes for the digital transmission of information.

be 23 distinct ways to stack copies of the Leech lattice together in 25-dimensional space: there are 23 different 25-dimensional laminated lattices. In 26 dimensions the laminated lattices number at least 75,000.

In dimensions higher than 25 much less is known. Hermann Minkowski demonstrated in 1905 that there are lattice packings in any dimension n whose density is greater than 2^{-n} . His argument, like others I have mentioned, is nonconstructive. In low dimensions the estimate is rather crude: when n is equal to 24, for example, it states merely that there exist packings whose density exceeds about 6×10^{-8} , whereas the Leech lattice is known to be more than 32,000 times denser. In the other direction Blichfeldt showed in 1914 that for arbitrarily large numbers n the density cannot exceed about $2^{-.5n}$. In spite of much effort there was essentially no improvement on this bound until the 1970's. G. A. Kabatiansky, Levenshtein

and V. M. Sidel'nikov in Moscow then showed that the density in very high dimensions n cannot exceed about $2^{-.599n}$.

E. S. Barnes of the University of Adelaide, A. Bos of N. V. Philips' Gloeilampenfabrieken in Eindhoven, the Netherlands, Conway, Leech and I have constructed a number of explicit packings in high dimensions, but none of them is as dense as Minkowski's theorem promises. Recently Barnes and I constructed lattice packings from the Leech lattice in dimensions up to 100,000. The density of these packings is roughly $2^{-1.25n}$, which at first glance seems almost as dense as the bound in Minkowski's theorem. Indeed, in dimension 65,536 our packings are about $10^{40,000}$ times denser than any lattice packing previously known. Unfortunately the exponential form of the bound can obscure the fact that we still fall rather short of the promised goal. The theorem guarantees there remain packings to be discovered that are $10^{4,000}$ times denser than the ones we have found.